REMARKS

By the foregoing Amendments, claims 42, 52, 62 and 66 have been amended. These changes are not believed to introduce new matter and are believed to place the application in better form for appeal by materially reducing or simplifying the issues for appeal. Entry of the amendments is respectfully requested.

Based on the following Remarks, Applicants respectfully request that the Examiner reconsider all outstanding rejections, and that they be withdrawn.

Rejections under 35 U.S.C. §112, first paragraph

At paragraph 2 of the Office Action (Paper No. 13), the Examiner rejected claims 42-61 as containing subject matter that was not described in the specification in such a way as to reasonably convey to one skilled in the art that the inventors, at the time the application was filed, had possession of the claimed invention. In particular, the Examiner objected to the term "state machine."

As agreed upon during the Examiner interview of March 6, 2001, Applicants have accommodated the Examiner's rejection by removing reference to the term "state machine" in independent claims 42 and 52.

Rejections under 35 U.S.C. §103

At paragraphs 5-14 of the Office Action (Paper No. 10), the Examiner rejected claims 42-46, 48-56, and 58-61 as being unpatentable over U.S. Patent No. 5,790,130 to Gannett ("Gannett") in view of U.S. Patent No. 6,097,402 to Case et al. ("Case"). Additionally, at paragraphs 15 and 16 of the Office Action (Paper No. 10), the Examiner rejected dependent claims 47 and 57 as being unpatentable over Gannett in view of Case and U.S. Patent No. 5,926,187 to Kim.

In rejecting independent claims 42 and 52, the Examiner asserted that Gannett teaches all of the claimed features except for a DMA engine. The Examiner then relied on Case for the alleged teachings of a DMA engine. In combining the teachings of Gannett with the teachings of Case, the Examiner stated that "it would have been obvious . . . to have combined the teachings of Gannett and Case such that textures are downloaded using DMA for the purpose of directly transferring the texture data to the cache at a faster rate and without burdening the host CPU."

Applicants respectfully submit that Gannett does not teach what the Examiner alleges. Specifically, Applicants note that Gannett does not teach a "graphics accelerator including . . . a replacement control component that implements a replacement policy for [a] texture cache memory." In setting forth the rejection, the Examiner asserts that Gannett's texture interrupt managing software daemon 160 is applicable to Applicants' replacement control component. This assertion is in error.

Gannett's texture interrupt managing (TIM) software daemon is not applicable to Applicants' replacement control component because the TIM software daemon is an independent, stand-alone software process that runs on the processor of the host computer.[1] This software process implements a cache replacement algorithm and communicates with one or more graphics hardware devices (e.g., texture mapping board 12) over a bus.[2] Thus, as the alleged replacement control component of Gannett is part of the host computer and not a part of texture mapping board 12, Gannett does not teach the claimed "graphics accelerator including . . . a replacement control component."

For at least this reason, Applicants submit that the Examiner has not presented a *prima facie* case of obviousness. The rejection of independent claims 42 and 52 is therefore traversed. The rejection of claims 43-51 and 53-61, which depend either directly or indirectly from one of independent claims 42 and 52, is also traversed for at least those reasons stated above.

---

[1] *See* col. 8, lines 51-54 and col. 12, lines 20-23 of Gannett.
[2] *See* col. 8, line 66 to col. 11, line 67 and FIGS. 3, 3A, 3B, and 4 of Gannett.

Applicants further note that Gannett's implementation of a cache replacement algorithm as part of a TIM software daemon is illustrative of a fundamental difference between Gannett and Applicants' system. While, data transfers in Gannett's texture cache are controlled by the CPU,[3] data transfers in Applicants' system are controlled by a direct memory access engine.[4]

This fundamental difference has not been addressed by the Examiner in combining the teachings of Gannett and Case. Indeed, the Examiner has not provided any evidence to justify a combination of Gannett and Case. For at least this additional reason, Applicants submit that the Examiner has not presented a *prima facie* case of obviousness.

Rejections under 35 U.S.C. §102(e)

At paragraphs 18 and 19 of the Office Action (Paper No. 10), the Examiner rejected claims 62, 63, and 65-68 as being anticipated by Gannett. The Examiner's rejection appears to be based largely on Gannett's use of S and T bits in addressing texels within a cache. Applicants submit that Gannett does not teach Applicants' use of an "address that is formed by interleaving individual bit values of a coordinate in a first dimension with individual bit values of a coordinate in a second dimension."

Gannett's storage of texel data in cache memory is illustrated in FIG. 9. The addressing scheme of Gannett's cache memory appears to be based on conventional row and column addressing:

> The cache includes eight SDRAM chips labeled SD1-SD8 divided evenly among the four interleaves 204A-204D, with each interleave including two SDRAM chips. The two SDRAMs in each interleave share the following common lines: an address line (ADD), row and column address strobes (RAS and CAS), a write enable (WE), a clock enable (CKE) and a data input/output mask (DQM). . . . The SDRAM chips receive twenty address bits multiplexed on line ADD to decode the 1,048,576 8-bit words within each bank. As explained in detail below, a 6-bit block index and a

[3] *See*, e.g., col. 8, line 51 to col. 9, line 29 of Gannett.
[4] *See* pages 9-10 and FIG. 2 of Applicants' specification.

> 16-bit texel address are computed for each texel to be accessed from the cache. The block index indicates in which of the sixty-four blocks of data the texel is located and the texel address indicates the precise S,T coordinate address of the texel within the block. Eight S bits and eight T bits comprise the texel address assuming a square block of 256×256 texels. A cache address is a twenty-two bit word including the combination of the block index (six MSBs) and texel address (sixteen LSBs). The cache address indicates the precise location of the texel within the cache. During rendering, the tiler/boundary checker decodes the LSB S bit and LSB T bit of the texel address (i.e., the LSB S coordinate and the LSB T coordinate) to determine in which of the four interleaves of the cache the texel is stored. The remaining 20 greater address bits of the cache address are provided along the address line ADD to the two SDRAM chips within the appropriate interleave. Of the twenty address bits provided to the two SDRAMs, nine bits are used to select the column and eleven bits are used to select the row within the SDRAMs to access the texel data. As should be understood by those skilled in the art, the column and row address bits are separately latched into the SDRAMs on different cycles and the RAS and CAS strobes are used conventionally to access the data.[5]

As this excerpt demonstrates, Gannett addresses a 256x256 block of texel data using a texel address that includes eight S bits and eight T bits. There does not appear to be any indication by Gannett that the group of S bits and the group of T bits are combined in an interleaved fashion. On this point, the Examiner's rejection is silent as to how the texel address of Gannett applies to Applicants "address that is formed by interleaving individual bit values of a coordinate in a first dimension with individual bit values of a coordinate in a second dimension."

Applicants note that Gannett discusses interleaving in the context of SDRAM chips, not address bits. Specifically, Gannett discloses a cache memory that has four interleaves 204A, 204B, 204C and 204D. These interleaves enable the host computer to simultaneously access four adjacent texels in a MIP map that are stored in separate interleaves of memory. Applicants submit that the interleaving of SDRAMs in a cache memory is distinct from Applicants' "address that is formed by interleaving individual bit

[5] *See* col. 23, lines 5-52 of Gannett.

values of a coordinate in a first dimension with individual bit values of a coordinate in a second dimension."

The interleaving of address bits is described at pages 15-17 and FIGS. 5 and 6 of Applicants' specification. In the illustrated embodiment, the texels are stored in linear cache lines that are accessible using an address formed by interleaving individual bit values of texture coordinates in the 'u' and 'v' dimensions. In general, this storage process maintains the positional relationship of texels in the texture cache, thereby providing rapid access to texel values.

For at least these reasons, Applicants submit that Gannett does not teach all of the limitations of independent claims 62 and 66. Accordingly, the teachings of Gannett cannot anticipate Applicants' independent claims 62 and 66, as well as dependent claims 63, 65, 67, and 68. The Examiner's rejection is therefore traversed.

At paragraphs 22 and 23 of the Office Action (Paper No. 10), the Examiner also rejected dependent claims 64 and 69 as being unpatentable over Gannett in view of U.S. Patent No. 5,945,997 to Zhao et al. This rejection is based on the anticipation rejection discussed above. For at least those reasons discussed above, Applicants submit that the rejection of claims 64 and 69 is also traversed.

Conclusion

All of the stated grounds of rejection have been properly traversed, accommodated, or rendered moot. Applicants therefore respectfully request that the Examiner reconsider all presently outstanding rejections, and that they be withdrawn.

Dated: 3/12/01

Respectfully submitted,
COOLEY GODWARD LLP

By: _____
Duane S. Kobayashi
Reg. No. 41,122

Cooley Godward LLP
Attn: Patent Group
One Freedom Square - Reston Town Center
11951 Freedom Drive
Reston, VA 20190-5601
Tel: (703) 456-8000
Fax: (703) 456-8100
DSK:bjp

Appendix: Amended Claims in Marked-Up Form

42. (Amended) A computer system, comprising:

a bus;

a central processing unit coupled to said bus; and

a graphics accelerator coupled to said bus, said graphics accelerator including a texture cache system [implemented as a state machine], said texture cache system including a texture cache memory that stores texels to be used by a texel value generating circuit, a replacement control component that implements a replacement policy for said texture cache memory, and a direct memory access engine that retrieves texel data from memory.

52. (Amended) A texture mapping method using a texture cache system [implemented as a state machine], said texture cache system including a texture cache memory, a replacement control component, and a direct memory access engine, comprising:

(a) retrieving texels from memory via the direct memory access engine;

(b) storing said retrieved texels in the texture cache memory in accordance with a replacement policy that is determined by the replacement control component; and

(c) rendering a polygon using texels that are stored in the texture cache memory.

62. (Amended) A computer system, comprising:

a memory; and

a memory control that stores two-dimensional data in said memory, wherein said data is stored in said memory using an address that is formed by interleaving individual

[bits of] bit values of a coordinate in a first dimension with individual bit values of a coordinate in a second dimension.

66. (Amended) A texture caching method, comprising:

(a) identifying a set of two-dimensional data that is to be transferred into memory; and

(b) storing said set [if] of two-dimensional data in memory using an address that is formed by interleaving individual [bits of] bit values of a coordinate in a first dimension with individual bit values of a coordinate in a second dimension.